Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 44 (2015) 588 – 598

Procedia Computer Science

2015 Conference on Systems Engineering Research

A network-based approach to organizational culture and learning in system safety

Mark S. Avnet[a,]*

[a]Texas A&M University, 4075 Emerging Technologies Building, 3131 TAMU, College Station, TX 77843-3131, USA

Abstract

While it is now generally agreed that system safety cannot be adequately addressed using technical analysis alone, an approach to modeling the organizational issues associated with safety is still needed. This paper offers an analytical approach to assessing the complex relationships among organizational culture and safety practices and outcomes. The paper argues that, in principle, organizational culture can be represented as a network of shared mental models (SMMs). While it would be impractical to construct a network that fully captures an organization's culture, the approach can be used to model particular dimensions of culture. Thus, a network of SMMs is a meaningful representation of safety culture to the extent that the data effectively capture shared knowledge about system safety. Similarly, organizational learning can be quantified as the evolution of that network's structure over time. The goal of the research is to develop a quantitative methodology for analyzing the relationship of organizational culture and learning to safety performance. The research is built on a collaborative effort between academia and industry focused on improving process safety in the oil and gas industry, but it can be applied to safety-related problems across organizations. The results are expected to have implications for training, professional development, safety protocols, and methods for measuring and managing safety practices in the development and operation of complex engineered systems.

Peer-review under responsibility of the Stevens Institute of Technology.

Keywords: system safety; process safety; safety culture; organizational culture; organizational learning; shared mental models; social network analysis; operational performance

1. Introduction

System safety has emerged as a critical issue in recent years as large-scale engineered systems have become larger and more complex. While much of this is driven by technical complexity (larger numbers of components with

* Corresponding author. Tel.: +1-646-286-4739; fax:+1-979-458-4299.
E-mail address: avnet@tamu.edu

Peer-review under responsibility of the Stevens Institute of Technology.
doi:10.1016/j.procs.2015.03.061

Nomenclature

SMM	Shared mental model
LNG	Liquefied natural gas
PSM	Process Safety Management
OSHA	Occupational Safety and Health Administration
$S_{x,y}$	Mental model sharedness between persons *x* and *y*
MKOPSC	Mary Kay O'Connor Process Safety Center

more and more interactions among them), some of the most difficult safety-related problems have more to do with organizational systems and structures than with specific technical issues. As Reason argues, these "latent failures" are defined by their presence "within the system well before the onset of a recognizable accident sequence."[1] Leveson points out that the proximate technical cause of an accident is often only a symptom of a broader and more systemic problem.[2] Even if the hazard seems to be of a purely technical nature without an apparent organizational cause at the time of the incident, the existence of the technical problem often can be attributed to "inadequate control over the [engineering design and] development process" rather than over operations.[2] So, in such a case, the organizational dimension is still relevant but must be considered further upstream during engineering design.

The goal of this paper is to introduce an emerging research program focused on a systems-oriented approach to organizational culture and its relationship to safety practices and outcomes. In this section, a brief review of the literature on organizational culture and system safety is presented, and the case for a systems view of safety culture is made. Then, a network-based approach for modeling organizational culture and learning is introduced. Following that, some of the details of the methodology are presented. Finally, the application of this research to process safety in the oil and gas industry is discussed.

This section reviews the interdisciplinary literature that forms the basis for the proposed systems approach to safety culture. The first subsection discusses how the term "culture" has been defined and used in the literature and presents an argument for applying the concept of culture to system safety. The second subsection then discusses the motivation for analyzing safety culture from a systems perspective.

1.1. Organizational culture and safety performance

Schein defines culture as "a pattern of shared basic assumptions learned by a group as it solved its problems of external adaptation and internal integration, which has worked well enough to be considered valid and, therefore, to be taught to new members as the correct way to perceive, think, and feel in relation to those problems."[3] Still, culture is both "broad and deep" and includes many different dimensions of an organization and its people. For this reason, any study of organizational culture should focus on specific elements of culture with the goal of addressing a particular observed phenomenon.[3] While Schein's definition is the one being used here, the scope of the present research is limited to safety culture, i.e., the particular elements of culture that are related to system safety. In terms of Schein's definition, the "problems" solved and the resulting "way to perceive, think, and feel" are those that have direct or indirect implications for safety practices and outcomes.

Most researchers and practitioners agree that culture, however it is defined, plays an important role in virtually every aspect of an organization's performance. Based on research and professional experience in process safety, Mannan et al. developed a list of 10 attributes for creating a "Best-in-Class safety culture." Although the authors point out that an organization does not need to possess all of these attributes to achieve excellence in safety performance, organizations that do have exemplary safety records tend to demonstrate some subset of these attributes.[4] In a study of 500 organizations conducted over a 10-year period, Keller and Price examined the relationship between sustained operational performance and an aspect of culture that they call "organizational health," a metric based on survey data assessing 37 specific management practices.[5] Fig. 1 shows the authors' analysis of the relationship between performance and health among several refineries of a particular oil company. These results demonstrate a positive correlation between an organization's focus on health-related (i.e., cultural)

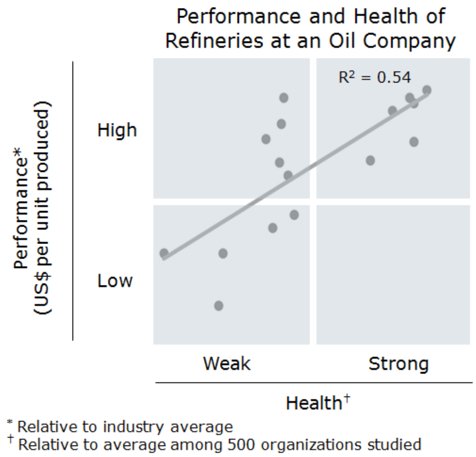

Fig. 1. The relationship of organizational health to performance at multiple refineries of one oil company.[5]

dimensions and operational performance. While this result is compelling, the analysis requires data from exhaustive (and very long) surveys. Bergman et al., on the other hand, have shown that similar surveys, focused in this case on "safety climate" in particular, may have a relatively short "shelf life." In other words, these types of surveys might be valid for predicting safety-related incidents only within a window of a few months after the date of the survey.[6] Given this outcome, the surveys used in the present research are designed to be short and targeted so that they can be distributed throughout the organization on a regular basis.

This subsection has discussed the relationship of culture to operational and safety performance. On the surface, it can appear as though these two priorities are in conflict. Indeed, as will be discussed in the next subsection, an overemphasis on operational performance can have a negative impact on safety in the short term.[1] Similarly, taking measures to improve safety can slow operations on a day-to-day basis. On the other hand, such safety interventions can improve operational outcomes in the longer term. When the organizational, economic, and political realities of a major safety incident are taken into account, the incident necessarily has a negative impact on performance. In other words, safety and productivity are complementary objectives.[1] The relationship between operational and safety performance, however, is not straightforward. These two priorities are linked in complex ways that involve many other aspects of the overall system. Thus, safety culture is best examined from a systems perspective.

1.2. Motivation for the systems perspective to safety

The idea of approaching safety – particularly in the context of complex engineered systems – from a holistic systems perspective that focuses on not only on technical causes of accidents but also on organizational factors is not a new one. The U.S. Navy's SUBSAFE program for nuclear submarines, which was established in 1963 after the loss of the USS Thresher, is guided by "a basic set of risk management principles, both technical and cultural."[2] The program suffered one early setback in 1968 when a vessel that was not SUBSAFE certified was lost. Following that incident, the Navy recommitted to SUBSAFE, and the program has been an unmitigated success since that time. The strength of the SUBSAFE program derives from the "constructive tension" that exists among the Platform Program Manager, an Independent Safety and Quality Assurance Authority, and an Independent Technical Authority. Supported by a culture of accountability and commitment from senior Navy leadership, these three entities are able to collectively maintain an effective and safe nuclear submarine program.[2]

Despite the success of SUBSAFE, the basic approach of emphasizing cultural aspects and ensuring accountability has not been widely emulated. Following the 1986 loss of the Space Shuttle *Challenger*, for example, the Rogers Commission recommended the formation of an independent oversight entity that would be held accountable for system safety but not for cost and schedule.[7] Although NASA established an independent safety office at Headquarters in Washington, DC in response to the report, the office did not have the enforcement authority to be effective. Furthermore, during the era of "faster, better, cheaper" under Administrator Daniel Goldin in the 1990s, the responsibility for this oversight was pushed back to the human spaceflight program. This act compromised the independence of the safety office and ultimately contributed to the *Columbia* tragedy in 2003.[8]

This climate of operational pressures, inadequate oversight, and failure to learn from past mistakes is not limited to NASA. The oil and gas industry, for example, sustained its own tragic loss with the explosion at the BP refinery in Texas City, Texas on March 23, 2005, which resulted in 15 deaths and 180 injuries. Among its organizational findings, the U.S. Chemical Safety and Hazard Investigation Board cited cost and production pressures, a lack of effective safety oversight, and the absence of "a reporting and learning culture" as key contributing factors to the accident.[9] Similarly, in a separate analysis triggered by the accident but more broadly focused on overall process safety across BP's five U.S. refineries, the Baker Panel criticized BP for its inadequate process safety leadership and decentralized safety management system. Although clear differences exist across the five refineries, the panel identified "significant process safety culture issues" at each of them.[10]

According to Leveson, the weaknesses inherent in the safety programs at NASA, BP, and countless other organizations are rooted in the fundamental approach to the problem. Although the traditional view of safety engineering as a reliability problem might have been adequate in the past, modern engineered systems have reached a level of complexity at which it is not possible to effectively model safety in terms of the combined probabilities of individual component failures. In reality, most accidents cannot be understood or prevented by focusing on component reliability. Instead, Leveson applies systems engineering principles to argue that the true cause of most accidents lies in the interactions among components. Moreover, these interactions are not limited to technical components but involve complex feedback and control among both technical and organizational elements. Based on this view, Leveson suggests that the entire field of safety engineering needs to be rethought. Safety is not a reliability problem; it is a control problem.[2]

To understand this view of safety as a control problem, consider the interactions depicted in Fig. 2. According to this model, a controller (human or machine) controls a process by initiating certain actions that are based on feedback from the controlled process. The rules that determine how the controller translates that feedback into an action are based on the controller's process model. For a computer controller, the process model is typically an algorithm. For a human controller, the process model is referred to as the person's mental model of the system or process.[2] This control-based perspective can apply to technical hazards in which the problem comes not from a component failure but from a flaw in the complex interactions among hardware and/or software components. Beyond that, however, this model also can often capture the organizational issues related to an accident even more effectively than it does the technical issues. For example, Leveson points out that both the *Challenger* and *Columbia* accidents "involved inadequate controls in the launch-decision process."[2]

At first glance, it may appear that multiple human controllers' process models for an engineered system should be roughly the same, but this is seldom the case. Consider how a designer's mental model might differ from the functioning of the actual system. The designer's understanding of the system is generally based on an idealized state derived from equations and software models, but the real-world conditions in which the system exists differ from

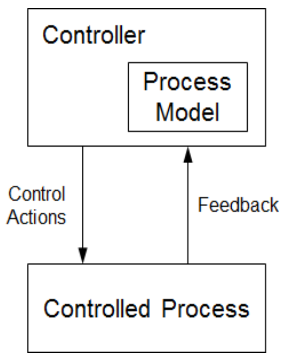

Fig. 2. Leveson's control model of system safety.[2]

that ideal state in important ways. The operator, in turn, develops his/her own mental model based on some combination of the designer's specifications and his/her operational experience with the actual system.[2] These differences in mental models can propagate across the entire system, resulting in a condition in which every individual involved – whether operator, designer, or manager – has a somewhat different mental model of the system. The next section presents an approach for studying the extent of similarities and differences across the mental models of all individuals involved in the system.

2. Shared knowledge, culture, and learning

Senge defines a learning organization as one "where people continually expand their capacity to create the results they truly desire, where new and expansive patterns of thinking are nurtured, where collective aspiration is set free, and where people are continually learning how to learn together."[11] According to Senge, becoming a learning organization requires five "disciplines": personal mastery, mental models, shared vision, team learning, and systems thinking. Systems thinking, Senge argues, is essential to ensuring that the other four disciplines work together as a coherent whole.[11] The remainder of this paper attempts to develop a formal and rigorous methodology for measuring an organization's approach to system safety against this goal.

The purpose of this section is to present an approach for analyzing shared knowledge in an organization using social network analysis. First, the literature on shared mental models is reviewed. Then, the basic principles of social network analysis are explained, and the rationale for representing organizational culture as a social network of shared mental models is described. Finally, some background literature on organizational learning is provided, and an argument for representing organizational learning as the time evolution of shared knowledge networks is presented.

2.1. Shared mental models

The term "mental model" is generally used to describe the way in which an individual perceives his or her environment. Some authors refer to a mental model as a catch-all for any knowledge about a given environment while others use the term only to describe organized knowledge that helps one to "understand phenomena, make inferences, and experience events by proxy."[12] Rouse and Morris define mental models more specifically as "mechanisms whereby humans are able to generate descriptions of system purpose and form, explanations of system functioning and observed system states, and predictions of future system states."[13] Based on this definition, two people can be said to hold a shared mental model (SMM) if they utilize mechanisms that lead to similar descriptions, explanations, and predictions of the system.

To ensure that the concept is useful across a variety of contexts, a distinction is often drawn among different types of SMMs based on their underlying content. Cannon-Bowers et al. classify mental models into two categories: task mental models (those that facilitate accomplishing a task) and team mental models (those that allow each individual to work effectively as a member of the team).[14] In two related studies, Mathieu et al. drew an empirical distinction between team and task mental models in two-person teams, or dyads. To operationalize SMMs, they computed a score based on the participants' perceptions of relationships among several task and team attributes.[15,16]

Kameda et al. conducted a study to measure the effect of "cognitive centrality" in teams of three, or triads. To do this, they constructed what they call a sociocognitive network and proposed a measure of cognitive centrality in the network.[17] Lim and Klein took this idea of examining shared cognition in larger teams a step further.[18] Based on a field study of 71 seven- to eight-person air combat teams, they devised a means of measuring shared knowledge for the entire team. Rather than looking only at shared knowledge in dyads as is done in much of the literature, they computed the average level of sharedness among all possible pairs of team members. The authors measured mental model accuracy by computing the same metric for each team member against experts' responses about the task.[18]

The above study made an important contribution by extending shared mental models to larger teams, but the approach was still a simple mean-based aggregation of pair-wise shared mental models. It did not provide a means of analyzing patterns of interactions among mental models. Langan-Fox et al. describe such patterns as "a synergistic functional aggregation of the [team's] mental functioning representing similarity, overlap, and complementarity."[19] Similarly, Klimoski and Mohammed refer to this phenomenon as "an emergent characteristic of the group, which is more than just the sum of individual models."[12] Building on this view of shared knowledge as emerging from SMMs among individuals, Avnet and Weigel developed a methodology for measuring shared knowledge based on the notion that teams are themselves complex systems in which people are the components.[20] While Lim and Klein measured averages of shared mental models across the team,[18] Avnet and Weigel extended the approach by examining complex patterns of relationships across all pair-wise SMMs in the team using social network analysis.[20] The next subsection explains the concept of social network analysis and then discusses how a network of shared knowledge can provide a system-level representation of safety culture in an organization.

2.2. Social networks and the structure of shared knowledge

Network analysis refers to a set of methods and techniques used to understand the global properties among a group of interacting entities in a complex system. A network consists of nodes, which represent individual entities, and edges (or arcs), which connect the nodes according to some type of relationship or interaction. A network is directed if its edges represent directional flow from one node to another. Otherwise, the network is undirected. Social networks are those used to analyze communication patterns or relationships among people in organizations.[21,22] In this type of network, a node usually represents a person, and an edge represents a measurable communication or a relationship between two people. Because of its emphasis on global patterns of interactions among large numbers of entities, social network analysis is scalable to networks of any size (limited pragmatically, of course, by computational capacity).

The purpose of studying networks is to understand broad properties of the overall structure of a system of interacting entities. For this reason, network analysis is often also called structural analysis.[21] Avnet and Weigel applied this perspective to shared knowledge in engineering design teams.[20] In their undirected networks of SMMs, each node represents a team member and each edge the shared mental model between a pair of team members. Two specific examples of these shared knowledge networks are shown in Fig. 3. The pattern of edges in such a network constitutes the structure of shared knowledge in the team. The authors therefore call this technique the structural approach to shared knowledge.[20]

Given the scalability of social network analysis, the structural approach is applicable to teams of any size and to large organizations. In this context, the structure of the shared mental model network is a representation of the "pattern of shared basic assumptions" that make up culture according to Schein's definition.[3] The limitation to the network's usefulness as a representation of culture then lies in the content of the shared mental models represented by each edge. To the extent that the survey questions used effectively capture shared knowledge about safety, it can be said that such a network is a quantitative representation of the organization's safety culture.

In real-world organizations, it is unlikely that the network representation of shared knowledge will stay constant over time. At most, the network model describes a particular aspect of organizational culture – in this case safety culture – at a snapshot in time. In other words, it is a static view of safety culture. In the next subsection, a dynamic view of this type of network is described, and an argument is presented for using the time evolution of the network as a measure of organizational learning.

2.3. Shared knowledge and organizational learning

According to Leveson's description of a system's safety control structure, the process model held by an individual human controller is that person's mental model.[2] Thus, the shared mental model held by a pair of such individuals is a representation or a result of the interactions between their process models. If a problem arises in the way that control actions and feedback are transmitted, these mental models could become inaccurate representations of the system. This condition would be exacerbated as the process models interact with and influence each other indirectly through other feedback and controls within the system. Leveson argues that most instances of "human error" can be traced to system-level issues that result in a mismatch between the individual's process model and the actual system.[2]

Given the set of complex feedback loops that Leveson describes in many modern systems,[2] it can be said that the process model of one controller can influence several aspects of the controlled system, which in turn affects other human controllers' process models. In the case of a "mismatch" as described above, this set of complex interactions

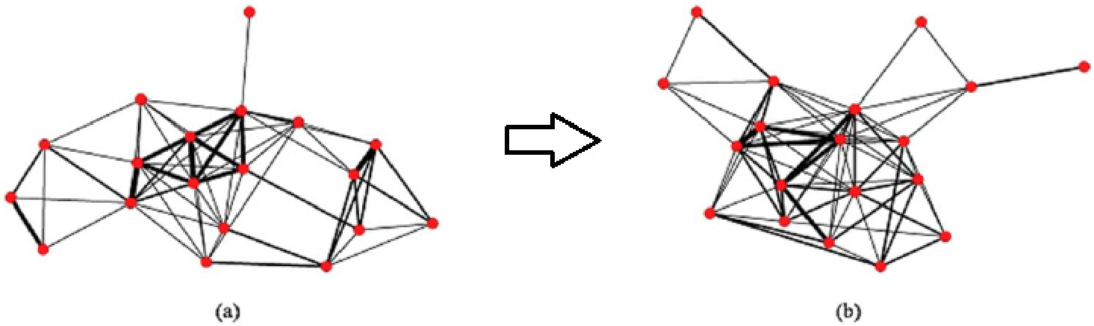

Fig. 3. Structure of shared knowledge in an engineering design team at two points in time: (a) before the work was completed and (b) after the work was completed.[20]

among mental models constitutes group think, the potential drawback of shared mental model development. On the other hand, when individuals' process models are appropriately aligned and converge over time toward a true representation of the actual system (i.e., they develop "correct" mental models), this type of change can be viewed as learning in the organization.

In the literature, organizational learning is defined as "a change in the organization's knowledge that occurs as a function of experience."[23] For the purposes of the present research, the "organization's knowledge" is operationalized as the structure of the network of shared mental models. Thus, according to the definition cited above, organizational learning can be viewed as an overall change in the structure of shared knowledge over time. This view of learning is consistent with observations that "organizational learning is not simply the sum of each member's learning."[24] That is, organizational learning is an emergent property that depends not only on individual learning but also on the pattern of interactions among the learning of all individuals.

In their study of systems engineering teams working on early conceptual design of scientific spacecraft, Avnet and Weigel developed a metric for change in shared knowledge based on an edge-by-edge correlation of a shared knowledge network at different points in time. The evolution of one such network is shown graphically in Fig. 3. Based on this analysis, the authors demonstrated a statistically significant correlation between change in shared knowledge and each of several technical characteristics of the system. Two examples of these relationships are shown in Fig. 4. In addition, the authors also found a correlation between change in shared knowledge and team coordination. Because of the difficulty in assessing quality of outcomes in early-stage conceptual design, the authors were not able to draw conclusions about the relationship of change in shared knowledge to performance.[20] Several studies in other contexts, however, have provided evidence that organizational learning is related to improved performance.[25,26]

This section has presented the rationale for studying safety culture and organizational learning using networks of shared mental models about factors that influence safety in the organization. In the next section, the specific methods used to construct and analyze these networks are presented.

3. Methods

This section describes a systems-oriented approach for analyzing organizational culture and learning. The approach is built on the methods developed by Avnet and Weigel, and the reader is referred to that study for a more complete description of its methods.[20] The present research differs from those methods in three important ways. First, this study uses survey questions on factors that influence safety rather than questions on engineering design drivers. Second, the surveys include additional questions that address cultural as well as technical and operational issues. Finally, this research takes advantage of the inherent scalability of network analysis by studying shared knowledge networks in large organizations rather than in teams. The author has partnered with Sentis, a global safety culture consultancy, to pilot the survey tools in a real-world context within the offshore oil and gas industry.

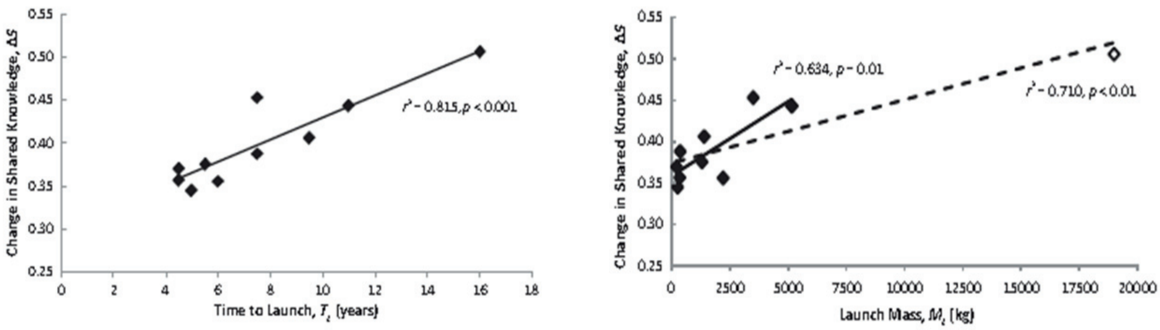

Fig. 4. Relationship of change in shared knowledge among members of engineering design teams to two technical attributes of the designed system: development time and system mass. Adapted from Avnet and Weigel.[20]

Sentis has piloted the survey in a specific systems engineering context, the construction of an undersea liquefied natural gas (LNG) extraction and refining platform in the Oceania region. The survey respondents are employees of a Sentis client contracted to complete this work for one of the world's leading oil and gas manufacturers.

At the time of this paper's writing, data collection has only recently begun. As such, this section discusses the data currently being collected and the planned analysis. The section is organized as follows. In the first subsection, the approach to measuring shared mental models and constructing shared knowledge networks in the context of system safety is described. In the second subsection, the method for measuring structural changes in shared knowledge networks over time is briefly reviewed. Finally, in the last subsection, the intended analysis to determine the relationship of learning to safety practices and outcomes is briefly discussed.

3.1. Constructing a network of shared knowledge

To define their metric for shared knowledge, Avnet and Weigel conducted an extensive literature review of prior studies measuring shared mental models in operational teams. Based on this work and on the factors that distinguish engineering design from operations, they created a survey asking design team members to check boxes corresponding to aspects of the system that most significantly drive the design work.[20] In early stage conceptual design, relationships among factors are generally open-ended or are not yet fully articulated. Thus, it would be difficult for a respondent to provide meaningful granularity in his/her responses to these questions. For this reason, simple check boxes were used instead of the Likert scales employed in most of the prior studies that focused on operations rather than design.[15,16,18] Data collection for the present research is similar to the approach taken by Avnet and Weigel but has been tailored to safety culture in the oil and gas industry. Thus, the survey is not based on system design drivers but instead asks respondents to rate various technical and organizational factors that influence safety. Furthermore, because many of the questions in this context necessarily involve the respondents' degree of agreement with various statements about each of the safety factors examined, a Likert scale is more appropriate than it was in the engineering design context. Depending on the organizational context, answer choices can include such issues as codification of protocols, adherence to procedures, tribal knowledge, management emphasis on safety and on production, risk tolerance in the group and in the organization, personal dynamics among operators, union rules, government regulations, safety-related features of equipment, and other specifics of the organization and/or technology that vary across companies and processes. For the Sentis pilot study, the factors chosen were based on three types of sources: the 14 elements of process safety management (PSM) as established by the U.S. Occupational Safety and Health Administration (OSHA)[27] and equivalent documents used in other countries, interviews with Sentis experts on process safety in the oil and gas industry, and documentation of specific practices and procedures used by the client organization. A portion of the survey currently in use is shown in Fig. 5.

The survey data will be used to calculate a metric of mental model sharedness, $S_{x,y}$, for each pair of respondents x and y. In the engineering design context, the metric was carefully chosen based on specifics of the design context and was validated, as shown in Fig. 4, against technical characteristics of the engineered system being designed.[20]

	To what extent are each of the listed factors important for achieving safe production?	Don't know	Not at all	Little extent	Some extent	Great extent	Very great extent
17	Clearly documented procedures and protocols for all job functions	☐	☐	☐	☐	☐	☐
26	Cross-training (traininig in other teams' roles and duties)	☐	☐	☐	☐	☐	☐
27	Emphasis on execution	☐	☐	☐	☐	☐	☐
28	Emphasis on problem solving	☐	☐	☐	☐	☐	☐
29	Pre-starts and after action reviews	☐	☐	☐	☐	☐	☐
30	Tracing incidents to individuals that make critical mistakes	☐	☐	☐	☐	☐	☐
39	Sharing information about safety problems with others	☐	☐	☐	☐	☐	☐
44	A positive relationship with my direct leader	☐	☐	☐	☐	☐	☐
47	Within-team cooperation and communication	☐	☐	☐	☐	☐	☐
48	Onshore management safety commitment	☐	☐	☐	☐	☐	☐

Fig. 5. Sample survey questions distributed to the construction crew for a liquefied natural gas (LNG) extraction and refining platform.

In the case of system safety, a similar process will be used to ensure that the chosen metric captures the important features related to safety culture and will be validated by relating its values under various conditions to known features of a given organizational context (e.g., number of employees, production volume, complexity of technical processes). A network of shared knowledge will then be constructed with $S_{x,y}$ as the edge between each pair of nodes x and y, and various metrics characterizing the network will be calculated. In engineering design, Avnet and Weigel measured relationships of the evolution of the shared knowledge network over time to various features of the technical design work, but they did not identify any structural features of each static network that related to the technical attributes examined.[20] This, however, will not necessarily be the case in the context of system safety. Thus, centrality measures will be used to determine the relative importance of knowledge held by each individual at a point in time. Although it is widely acknowledged that leadership support is critical to implementing and sustaining change in an organization,[3,5] other members of the organization could occupy central positions in the network based on their personalities and/or their own approaches to safety. Krackhardt and Hanson have analyzed informal networks in organizations, classifying such networks as advice, trust, and communication.[28] The present research will include an investigation of the possibility of shared knowledge networks serving as a fourth type of informal network – the knowledge network.

3.2. Measuring organizational learning

Avnet and Weigel calculated an edge-by-edge correlation of each shared knowledge network at different points in time to determine the extent to which shared knowledge remained stable over time, and a simple transformation of this similarity measure resulted in a metric for change in shared knowledge.[20] As shown in Fig. 3, the authors constructed a network of shared knowledge before and after each design session. Because the sessions each lasted for only five days, it was not possible to measure precisely how shared knowledge evolved over time.

Similarly, in the Sentis pilot study, the crew has been deployed for a well-defined and relatively short mission, which does not allow for multiple time points to be used. Thus, this pilot study includes only surveys at the start and at the end of the work, as was done in the engineering design context.[20] As the research develops, however, the shared mental model surveys will be distributed regularly. Depending on the specifics of the organizations studied, distribution could be weekly, semi-weekly, or monthly. Survey frequencies will likely be adjusted based on the willingness of participants and on the time scales over which meaningful changes seem to occur early in the study. Using these time series data, shared knowledge networks will be constructed at multiple points in time. This will make it possible to plot a trajectory of changes in shared knowledge over time. The shape of the resulting curve can then be used to understand how safety-related learning takes place in an organization.

3.3. Relationship to safety practices and outcomes

This research program will involve measuring technical performance, adherence to safety practices, and safety outcomes in the organizations studied. For the pilot study, data on safety performance outcomes will be obtained from the client at the conclusion of the platform construction project. In addition, Sentis has conducted structured interviews with more than 20 crew members and client personnel to assess safety practices, and the analysis of the interview data is ongoing at the time of this paper's writing. After the data have been collected, the findings about organizational culture and learning determined from the SMM networks will be related to safety performance metrics. It is not yet clear what the exact performance metrics will be for Sentis' client or for other organizations studied. The specific metrics used will undoubtedly be shaped, at least in part, by the particular organizations involved in the study. In the next section, the future directions of the research program are briefly discussed.

4. Conclusion and future directions

The nascent research program presented here began with a discussion among academics and practitioners from the oil and gas, chemical, and nuclear industries during a meeting of the Mary Kay O'Connor Process Safety Center (MKOPSC) at Texas A&M University. The MKOPSC serves as a hub for research and discussion on process safety, and it facilitates research partnerships among industry, government, and academia. While most of the work of the MKOPSC focuses on technical proximate causes, there is a growing awareness among its members that the path to improving process safety needs to include a strong emphasis on safety culture.

Since those initial discussions, members of the MKOPSC have met regularly to discuss a research program focused on the network approach to organizational culture and learning as presented in this paper. In addition to the partnership already established with Sentis to work with many of the company's clients, interest in the research is also growing steadily among major companies in the process industries. It remains to be seen how the research will progress in the coming months, but the author and members of the MKOPSC are encouraged by the potential of the Sentis pilot study and by the response that the research approach has generated so far from major oil and gas companies. Ultimately, the goal of this research is to demonstrate the complex organizational factors involved in system safety and to develop a systems-oriented methodology for measuring organizational culture and learning about safety. The outcomes of the work will have direct implications for organizational and technical issues relevant to the development of a variety of complex engineered systems ranging from the ocean-based platforms and oil refining facilities examined in this investigation to a host of other systems engineering contexts, including aerospace systems, transportation, and patient safety in healthcare systems. From a practical standpoint, this understanding can be used to develop more targeted incentive structures, communication mechanisms, standard procedures, training programs, and technical system features that will help to ensure the development and management of safer systems engineering organizations in the future.

Acknowledgements

The author would like to thank M. Sam Mannan and the staff of the Mary Kay O'Connor Process Safety Center for facilitating the early collaboration on which this research is based. The author also would like to thank Fred Infortunio for his valuable input to the initial shaping of the manuscript. Ray Mentzer and Dennis Parker provided important industry-specific feedback and examples, and David Broniatowski offered extensive commentary on an early draft. The author would like to extend a special thanks to Tristan Casey at Sentis for identifying the opportunity to incorporate this research into a client engagement, for helping to develop the survey tool, for collecting the data during the pilot study, and for editing a recent draft of the paper. In addition, the author would like to thank Autumn Krauss at Sentis for initiating the research partnership between Sentis and Texas A&M University, which promises to yield valuable impact in terms of both theory and practice as the program develops. Finally, the author would like to thank those that contributed to the previous study on which the methodology presented here is based: Annalisa Weigel, Christopher Magee, John Carroll, and Richard Binzel.

References

1. Reason J. The contribution of latent human failures to the breakdown of complex systems. *Phil Trans R Soc Lond B* 1990;**327**: 475-484.
2. Leveson N. *Engineering a safer world: systems thinking applied to safety*. Cambridge: MIT; 2011.
3. Schein E. *Organizational culture and leadership*. 4th ed. San Francisco: Jossey-Bass, 2010.
4. Mannan MS, Mentzer RA, Zhang J. Framework for creating a Best-in-Class safety culture. *J Loss Prevent Proc* 2013;**26(6)**:1423-1432.
5. Keller S, Price C. Beyond performance: how great organizations build ultimate competitive advantage. Hoboken: Wiley; 2011.
6. Bergman ME, Payne SC, Taylor AB, Beus JM. The shelf life of a safety climate assessment: how long until the relationship with safety-critical incidents expires? *J Bus Psychol* 2014;**29(4)**: 519-540.
7. Presidential Commission on the Space Shuttle Challenger. *Report to the president*. Washington: Government Printing Office; 1986.
8. Columbia Accident Investigation Board. *Report, vol. 1*. Washington: National Aeronautics and Space Administration and Government Printing Office; 2003.
9. U.S. Chemical Safety and Hazard Investigation Board. *Investigation report: refinery explosion and fire, BP*. Texas City: U.S. Chemical Safety and Hazard Investigation Board; 2007.
10. BP U.S. Refineries Independent Safety Review Panel. *The report of the BP U.S. Refineries Independent Safety Review Panel*. Texas City: BP U.S. Refineries Independent Safety Review Panel; 2007.
11. Senge PM. *The fifth discipline: the Art & practice of the learning organization*. New York: Doubleday; 2006.
12. Klimoski R, Mohammed S. Team mental model: construct or metaphor? *J Manage* 1994;**20(2)**:403-437.
13. Rouse WB, Morris NM. On looking into the black box: prospects and limits in the search for mental models. *Psychol Bull* 1986;**100(3)**:349-363.
14. Cannon-Bowers JA, Salas E, Converse SA. Shared mental models in team decision making. In: Castellan Jr NJ, editor. *Individual and group decision making*. Hillsdale: Erlbaum; 1993. p. 221-246.
15. Mathieu JE, Heffner TS, Goodwin GF, Salas E, Cannon-Bowers JA. The influence of shared mental models on team process and performance. *J Appl Psychol* 2000;**85(2)**:273-283.
16. Mathieu JE, Heffner TS, Goodwin GF, Cannon-Bowers JA, Salas E. Scaling the quality of teammates' mental models: equifinality and normative comparisons. *J Organ Behav* 2005;**26**:37-56.
17. Kameda T, Ohtsubo Y, Takezawa M. Centrality in sociocognitive networks and social influence: an illustration in a group decision-making context. *J Pers Soc Psychol* 1997;**73(2)**:296-309.
18. Lim B-C, Klein KJ. Team mental models and team performance: a field study of the effects of team mental model similarity and accuracy. *J Organ Behav* 2006;**27**:403-408.
19. Langan-Fox J, Anglim J, Wilson JR. Mental models, team mental models, and performance: process, development, and future directions. *Hum Factor Ergon Man* 2004;**14(4)**:331-352.
20. Avnet MS, Weigel AL. The structural approach to shared knowledge: an application to engineering design teams. *Hum Factors* 2013;**55(3)**:581-594.
21. Wasserman S, Faust K. *Social network analysis: methods and applications (structural analysis in the social sciences)*. Cambridge: Cambridge; 1999.
22. Newman MEJ. *Networks: an introduction*. Oxford: Oxford; 2010.
23. Argote L, Miron-Spektor E. Organizational learning: from experience to knowledge. *Organ Sci* 2011;**22(5)**:1123-1137.
24. Fiol CM, Lyles MA. Organizational learning. *Acad Manage Rev* 1985;**10(4)**:803-813.
25. Levitt B, March JG. Organizational learning. *Annu Rev Sociol* 1988;**14**:319-340.
26. Jimenez-Jimenez D, Sanz-Valle R. Innovation, organizational learning, and performance. *J Bus Res* 2011;**64(4)**:408-417.
27. Occupational Safety and Health Administration. *Process safety management: OSHA 3132*. Washington: U.S. Department of Labor; 2000.
28. Krackhardt D, Hanson JR. Informal networks: The company behind the chart. *Harvard Bus Rev* 1993;**73(2)**:104-111.